

ABSTRACT

A method and apparatus for ranked join indices includes a solution providing performance guarantees for top-k join queries over two relations, when preprocessing to construct a ranked join index for a specific join

5 condition is permitted. The concepts of ranking join indices presented herein are also applicable in the case of a single relation. In this case, the concepts herein provide a solution to the top-k selection problem with monotone linear functions, having guaranteed worst case search performance for the case of two ranked attributes and arbitrary preference vectors.